

07-20-00

A



IBM Docket No. RSW9-2000-0054-US1

JC675 U.S. PTO
09/6/9178

In the United States Patent and Trademark Office Patent Application Transmittal

Transmitted herewith for filing is the Patent Application of:

Inventors(s): Donald J. Boulia

For: Technique for Sending TCP Messages through HTTP Systems

Enclosed are

36 pages of specification, including 30 claims, plus 5 sheets of formal drawings.

X An assignment of the invention to International Business Machines Corporation, Armonk, New York 10504.

A certified copy of a/an application.

X Declaration and Power of Attorney.

PTO-1449 & references

X A return post card

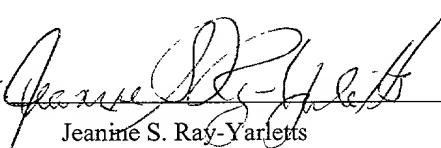
Other:

Filing Fee Calculation (For Other Than Small Entity)

Basic Fee:						\$690.00
Claims Fees:	Filed	Limit	Extra		Rate per Extra	
Total claims:	30	20	10		\$18.00	\$180.00
Independent claims:	4	3	1		\$78.00	\$78.00
Multiple Dependent Claim Presented					\$260.00	\$0.00
					Total	\$948.00

Please charge Deposit Account 09-0461 for the Total set forth above. The Commissioner is authorized to charge payment of any additional filing fees required under 37 CFR §1.16 and any patent application processing fees under 37 CFR §1.17 or to credit any overpayment to the identified account. A duplicate copy of this sheet is enclosed.

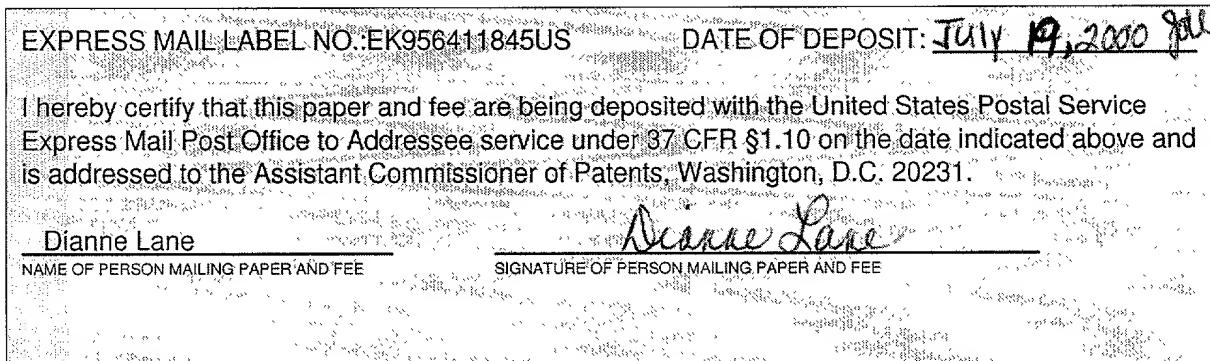
<i>Express Mail Certificate</i>	
Express Mail Label No. EK956411845US	
Date: July 19, 2000	
I hereby certify that I am depositing the papers identified above with the U.S. Postal Service "Express Mail Post Office to Address" service on the above date, addressed to the Commissioner of Patents and Trademarks, Washington, DC 20231	
<i>Dianne Lane</i>	
Dianne Lane	

BY 
 Jeanine S. Ray-Yarletts
 Attorney of Record Reg. No. 39,808

Date: July 19, 2000

IBM Corporation T81/062
 Intellectual Property Law
 PO Box 12195
 Res. Tri. Park, NC 27709

Telephone: 919-543-2541 FAX 919-254-4330



INVENTOR: Donald J. Boulia

Technique for Sending TCP Messages through HTTP Systems

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates to distributed computing, and deals more particularly with a method, system, and computer program product for sending TCP (Transmission Control Protocol) messages through a network using the HyperText Transfer Protocol (HTTP) and HTTP-based systems.

Description of the Related Art

Business and consumer use of distributed computing, also commonly referred to as

10 network computing, has gained tremendous popularity in recent years. In this computing model,

the data and/or programs to be used to perform a particular computing task typically reside on
(i.e. are “distributed” among) more than one computer, where these multiple computers are
connected by a network of some type. The Internet, and the part of the Internet known as the
World Wide Web (hereinafter, “Web”), are well-known examples of this type of environment
5 wherein the multiple computers are connected using a public network. Other types of network
environments in which distributed computing may be used include intranets, which are typically
private networks accessible to a restricted set of users (such as employees of a corporation), and
extranets (e.g., a corporate network which is accessible to other users than just the employees of
the company which owns and/or manages the network, such as the company’s business partners).

10 While businesses are eager to enable their business partners and customers to interact with
their computing resources for business-to-business and business-to-consumer electronic
commerce, those resources must be protected from unintended and malicious intrusion, and use of
computing resources must be made as efficient as possible. Many business enterprises have
therefore installed firewall systems and/or proxy systems in their enterprise networks. Firewall
systems are used to limit access to computing resources. A firewall system typically limits
15 incoming network traffic in an enterprise, as well as outgoing network traffic. Proxy systems act
as an intermediary between end users or clients connected to a network and a remote server
which is reachable through the network. Proxy systems are often associated with firewall
systems, and these functions may be combined in some server implementations. Firewall and
proxy functionality is typically implemented using a network traffic filtering approach. That is,
20 incoming traffic may be filtered to determine whether it originates from a source which is

permitted to access the enterprise's computing resources, and outgoing traffic may be filtered to determine whether it is destined for a remote location which meets certain predefined criteria. Firewall and proxy systems may also perform other services directed toward maximizing the efficiency of an enterprise's resources, such as caching and message logging.

5 The HyperText Transport Protocol (HTTP) is the communications protocol typically used for transmitting messages and data in the Web environment (and may be used in other networking environments as well). When a proxy is implemented for use with Web traffic or HTTP-based networks, it may be referred to as an "HTTP proxy".

10 Firewall and proxy systems, and the benefits they provide, are well known in the art.

15 However, these systems may introduce undesirable limitations for an enterprise as well. Many firewall systems place a limitation on the number of ports which may be opened, and severely restrict the number of ports which may be opened for non-HTTP traffic such as TCP (Transmission Control Protocol) traffic. TCP was developed long before HTTP, and many applications exist which were designed for use with TCP. End users within an enterprise have a continuing need to access these applications, which may be remotely located and therefore require access through an external network. Similarly, an enterprise may have many TCP-based applications to which it would like to provide access for external users. Because of the security benefits which result from limiting the number of open TCP connections on a firewall, simply providing more TCP ports is not a viable way to address this problem. Furthermore, when HTTP proxies are in place, these proxies often restrict the traffic which flows through them to HTTP-

20

based traffic, thereby adding further complications to those introduced by the firewall restrictions for transmitting TCP messages.

TCP is designed using a bi-directional computing model, wherein either party to a connection may initiate a message to the other. HTTP, on the other hand, is designed for communication using a transactional client/server model wherein requests are initiated by a client and responded to by a server: the protocol does not provide for server-initiated messages. The prior art therefore does not provide a technique for supporting TCP traffic as messages flowing through HTTP networks, nor for enabling TCP traffic to flow through an HTTP proxy of the type which has been discussed. At the same time, it is necessary to continue to provide end users with access to existing TCP-based applications. Rewriting or retrofitting the existing TCP-based application software to use a client/server model, and to use HTTP messages natively, would be extremely costly, time-consuming, and error-prone. This approach is therefore not a viable solution for most businesses.

Accordingly, what is needed is a technique for enabling TCP messages to be exchanged using HTTP networks and systems. This technique must not require change to existing client and server TCP-based software.

SUMMARY OF THE INVENTION

An object of the present invention is to provide a technique that facilitates end-user access to remote TCP-based applications.

Another object of the present invention is to provide a technique which enables TCP traffic to flow through firewalls without increasing the number of open TCP ports.

It is another object of the present invention to provide a technique which enables transmitting TCP traffic through HTTP proxy servers.

5

Other objects and advantages of the present invention will be set forth in part in the description and in the drawings which follow and, in part, will be obvious from the description or may be learned by practice of the invention.

10
11
12
13
14
15

15

To achieve the foregoing objects, and in accordance with the purpose of the invention as broadly described herein, the present invention provides a computer program product, a system, and a method for sending TCP messages through HTTP systems. This technique comprises: establishing a send channel from a first component on a client side of a network connection, through one or more HTTP-based systems, to a second component on a remote side of the network connection; establishing a receive channel from the first component, through the one or more HTTP-based systems, to the second component; establishing a first TCP connection from a client on the client side to the first component; establishing a second TCP connection from the second component to a target server on the remote side; transmitting client-initiated TCP requests from the client to the target server on the send channel; and transmitting server-initiated TCP requests from the target server to the client on the receive channel.

5

Transmitting the client-initiated TCP requests preferably further comprises: receiving a client-initiated TCP request from the client at the first component on the first TCP connection; packaging the received client-initiated TCP request in an HTTP POST request message; sending the HTTP POST request message to the second component on the network connection; receiving the sent HTTP POST request message at the second component; extracting the client-initiated TCP request from the received HTTP POST request message; and forwarding the extracted client-initiated TCP request to the target server on the second TCP connection. Transmitting the client-initiated TCP requests may further comprise acknowledging the HTTP POST request by sending an HTTP POST response from the second component to the first component on the network connection. The send channel is preferably established in response to receiving the client-initiated TCP request, and transmitting client-initiated TCP requests preferably further comprises receiving the HTTP POST response at the first component and then closing the send channel.

15

20

Transmitting the server-initiated TCP requests preferably further comprises: sending an HTTP GET request message from the first component to the second component on the network connection; receiving the sent HTTP GET request message at the second component; receiving a server-initiated TCP request from the target server at the second component on the second TCP connection; packaging the received server-initiated TCP request in an HTTP GET response message which acknowledges the received HTTP GET request message; sending the HTTP GET response message from the second component to the first component on the network connection; receiving the sent HTTP GET response message at the first component; extracting the server-

initiated TCP request from the received HTTP GET response message; and forwarding the extracted server-initiated TCP request to the client on the first TCP connection. Transmitting the server-initiated TCP requests may further comprise performing a read operation on the second TCP connection, responsive to receiving the sent HTTP GET request message and prior to receiving the server-initiated TCP request, and using the received server-initiated TCP request as a result of the read operation, thereby triggering the packaging of the received server-initiated TCP request in the HTTP GET response message. Transmitting the server-initiated TCP requests preferably also further comprises preparing to receive another server-initiated TCP request by triggering the sending of the HTTP GET request message from the first component to the second component, responsive to receiving the sent HTTP GET response message at the first component.

A Multi-Purpose Internet Mail Extensions (MIME) type of the HTTP POST request message, the HTTP POST response message, the HTTP GET request message, and/or the HTTP GET response message may be set to “binary/tcp”.

Rather than TCP and HTTP, the protocols involved may be an arbitrary bi-directional

15 protocol and uni-directional protocol, respectively.

The present invention will now be described with reference to the following drawings, in which like reference numbers denote the same element throughout.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram of a computer workstation environment in which the present invention may be practiced;

5 Figure 2 is a diagram of a networked computing environment in which the present

invention may be practiced;

Figure 3 provides a schematic illustration of the components which are used to implement

a preferred embodiment of the present invention; and

Figures 4A and 4B provide diagrams showing how the bi-directional message flow for

10 TCP-based applications is supported when messages are initiated by the client and by the server, respectively, according to a preferred embodiment of the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENT

Fig. 1 illustrates a representative workstation hardware environment in which the present invention may be practiced. The environment of Fig. 1 comprises a representative single user computer workstation 10, such as a personal computer, including related peripheral devices. The 15 workstation 10 includes a microprocessor 12 and a bus 14 employed to connect and enable communication between the microprocessor 12 and the components of the workstation 10 in accordance with known techniques. The workstation 10 typically includes a user interface adapter 16, which connects the microprocessor 12 via the bus 14 to one or more interface

devices, such as a keyboard 18, mouse 20, and/or other interface devices 22, which can be any user interface device, such as a touch sensitive screen, digitized entry pad, etc. The bus 14 also connects a display device 24, such as an LCD screen or monitor, to the microprocessor 12 via a display adapter 26. The bus 14 also connects the microprocessor 12 to memory 28 and long-term storage 30 which can include a hard drive, diskette drive, tape drive, etc.

5

The workstation 10 may communicate with other computers or networks of computers, for example via a communications channel or modem 32. Alternatively, the workstation 10 may communicate using a wireless interface at 32, such as a CDPD (cellular digital packet data) card. The workstation 10 may be associated with such other computers in a LAN or a wide area network (WAN), or the workstation 10 can be a client in a client/server arrangement with another computer, etc. All of these configurations, as well as the appropriate communications hardware and software, are known in the art.

10
15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95

15

The hardware environment of a server is well known in the art. The preferred embodiment of the present invention operates partly on the client side of a network connection and partly on the server side of a network connection. On the client side, the present invention is preferably embodied in a device such as workstation 10. On the server side of the network connection, the present invention is preferably embodied in a server such as a Web application server or a load-balancing server which may be installed as an intermediary between the network and an application server.

Fig. 2 illustrates a data processing network 40 in which the present invention may be practiced. The data processing network 40 may include a plurality of individual networks, such as wireless network 42 and network 44, each of which may include a plurality of individual workstations 10. Additionally, as those skilled in the art will appreciate, one or more LANs may be included (not shown), where a LAN may comprise a plurality of intelligent workstations coupled to a host processor.

Still referring to Fig. 2, the networks 42 and 44 may also include mainframe computers or servers, such as a gateway computer 46 or application server 47 (which may access a data repository 48). A gateway computer 46 serves as a point of entry into each network 44. The gateway 46 may be preferably coupled to another network 42 by means of a communications link 50a. The gateway 46 may also be directly coupled to one or more workstations 10 using a communications link 50b, 50c. The gateway computer 46 may be implemented utilizing an Enterprise Systems Architecture/370 available from IBM, an Enterprise Systems Architecture/390 computer, etc. Depending on the application, a midrange computer, such as an Application System/400 (also known as an AS/400) may be employed. (“Enterprise Systems Architecture/370” is a trademark of IBM; “Enterprise Systems Architecture/390”, “Application System/400”, and “AS/400” are registered trademarks of IBM.) These are merely representative types of computers with which the present invention may be used.

The gateway computer 46 may also be coupled 49 to a storage device (such as data repository 48). Further, the gateway 46 may be directly or indirectly coupled to one or more

workstations 10, and servers such as gateway 46 and application server 47 may be coupled to other servers such as server 43.

Those skilled in the art will appreciate that the gateway computer 46 may be located a great geographic distance from the network 42, and similarly, the workstations 10 may be located 5 a substantial distance from the networks 42 and 44. For example, the network 42 may be located in California, while the gateway 46 may be located in Texas, and one or more of the workstations 10 may be located in New York. The workstations 10 may connect to the wireless network 42 using a networking protocol such as the Transmission Control Protocol/Internet Protocol (“TCP/IP”) over a number of alternative connection media, such as cellular phone, radio frequency networks, satellite networks, etc. The wireless network 42 preferably connects to the gateway 46 using a network connection 50a such as TCP or UDP (User Datagram Protocol) over IP, X.25, Frame Relay, ISDN (Integrated Services Digital Network), PSTN (Public Switched Telephone Network), etc. The workstations 10 may alternatively connect directly to the gateway 46 using dial connections 50b or 50c. Further, the wireless network 42 and network 44 may connect to one or more other networks (not shown), in an analogous manner to that depicted in 15 Fig. 2. (Note that when the Host On-Demand product is used to access legacy host data, the IBM Systems Network Architecture and its protocols are also used within the network. It will be obvious to one of skill in the art how the configuration shown in Fig. 2 is augmented for this situation.)

20 Software programming code which embodies the present invention is typically accessed by

the microprocessor 12 (for example, of the workstation 10, server 43, gateway 46, and/or server 47) from long-term storage media 30 of some type, such as a CD-ROM drive or hard drive. The software programming code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users from the memory or storage of one computer system over a network of some type to other computer systems for use by users of such other systems. Alternatively, the programming code may be embodied in the memory 28, and accessed by the microprocessor 12 using the bus 14. The techniques and methods for embodying software programming code in memory, on physical media, and/or distributing software code via networks are well known and will not be further discussed herein.

10
15
20

The requesting and responding devices which make use of the present invention, when connected, may use a "wireline" connection or a "wireless" connection. Wireline connections are those that use physical media such as cables and telephone lines, whereas wireless connections use media such as satellite links, radio frequency waves, and infrared waves. Many connection techniques can be used with these various media, such as: using the computer's modem to establish a connection over a telephone line; using a LAN card such as Token Ring or Ethernet; using a cellular modem to establish a wireless connection; etc. The requesting computer may be any type of computer processor, including laptop, handheld or mobile computers; vehicle-mounted devices; desktop computers; mainframe computers; etc., having processing and communication capabilities. The responding computer, similarly, can be one of any number of different types of computer which have processing and communication capabilities. These

techniques are well known in the art, and the hardware devices and software which enable their
use are readily available. (The requesting computer is referred to alternatively herein as the
“requester” or “client” for ease of reference, although the requester may be a server machine
operating as a client for a particular request/response scenario, such as a server which initiates a
5 TCP request to another device. The responding computer is referred to alternatively herein as the
“requester” or the “server”, for ease of reference.)

In the preferred embodiment, the present invention is implemented in computer software.

The implementation of the software of the present invention may operate as one or more modules
(also referred to as code subroutines, or “objects” in object-oriented programming) on a client,
server, or intermediary device in a network. In the preferred embodiment, the JavaTM
10 programming language is used. (“Java” is a trademark of Sun Microsystems, Inc.) References
herein to use of the Java programming language, however, are for purposes of illustration and not
of limitation.

The present invention provides an applicant-independent technique for facilitating end-
15 user access to remote TCP-based applications. TCP traffic is packaged into the body of HTTP
messages, enabling the TCP traffic to be transmitted through HTTP systems as if it were native
HTTP traffic. This technique enables the TCP traffic to flow through firewall systems without
requiring access to external TCP ports on the firewall, and enables the TCP traffic to flow
20 through HTTP proxy servers as well. No change is required to existing end-user software, or to
the remote application software executing on a target host system or application server.

As an added benefit, packaging the TCP traffic into HTTP messages using the technique disclosed herein enables the TCP traffic to automatically take advantage of encryption which may be provided (using, for example, the HTTPS protocol, which is a security-enhanced version of HTTP) for the traffic flowing through particular firewalls or proxy servers.

5 The preferred embodiment of the present invention will now be discussed in more detail with reference to Figs. 3 and 4.

The components which may be used to implement the present invention are depicted in Fig. 3. The component which operates on the client side of a network connection is referred to herein as a “Web redirector”, and is shown as element 320. The component which operates on the server side is referred to herein as a “redirector servlet”, and is shown as element 360. (While the preferred embodiment describes the server-side component as a servlet, this is for purposes of illustration and not of limitation. An alternative implementation strategy may be substituted for the servlet model without deviating from the scope of the present invention.) The Web redirector 320 may execute on the same client device as the client’s remote access software, where that software is depicted in Fig. 3 as client 300. Alternatively, Web redirector 320 may execute on a separate device to which the client software makes a local connection. (In this latter case, Web redirector 320 may serve more than a single client 300.) The redirector servlet 360 is preferably located on a Web application server 350 or other device which connects to a target host or target application server (referred to hereinafter as a target server for ease of reference). The target server is shown in Fig. 3 as server 380.

The preferred embodiment is described herein with reference to an example in which the client and server use the Telnet protocol. The Telnet protocol is designed to provide remote terminal connection wherein a user of one computing device logs on to an application at a remote computing device as if the user is using a computing device which is locally-attached to the
5 remote device. Telnet is widely used for accessing legacy host applications and legacy host data, and provides for insulating the application from the details of the particular type of user computing device. While the preferred embodiment is discussed herein with reference to Telnet clients and Telnet servers, however, it is equally applicable to clients and servers which use other TCP-based protocols (such as File Transfer Protocol, or "FTP", Gopher, etc.). Furthermore, the present invention may be used advantageously for transporting messages for bi-directional
10 protocols other than TCP, including proprietary protocols, and may use underlying protocols other than HTTP for this transmission provided the alternative underlying protocol has features
15 which are semantically equivalent to those described herein for HTTP.

10
15

20

Client and server software which is written to communicate using TCP messages uses the concept of a connection between two sockets for exchanging data, where a socket is comprised of an address identifying one of the computers involved and a port number that identifies a particular process on that computer. The process identified by the port number is the process that will receive the incoming data for that socket. A socket is typically implemented as a queue, and the receiving process therefore retrieves data from its incoming socket queue. Port numbers may be temporarily assigned for use with a particular application, or may be permanently assigned such that they are, by convention, used with all implementations. Port numbers of the latter type are

referred to as “well-known port numbers”. The well-known port number for Telnet messages is port 23; the well-known port number for HTTP messages is port 80. These well-known port numbers are used herein for purposes of illustration and not of limitation.

According to the present invention, client software 300 continues to use port 23 for sending and receiving data on a TCP connection (shown in Fig. 3 as element 310), in the same manner it normally connects to a TCP connection, where the Web redirector 320 is now configured as the other party of that TCP connection. The software at server 380 also continues to use port 23 for sending and receiving TCP data, as in its normal operations, where the redirector servlet 360 is now at the other end of the TCP connection 370 to server 380. The client software and the software at a server therefore operates without change when using the present invention.

Web redirector 320 and redirector servlet 360 provide the bi-directional communication which is needed for exchanging TCP messages by establishing two channels, shown in Fig. 3 as a “receive” channel 330 and a “send” channel 340. The traffic which flows on channels 330 and 15 340 uses HTTP messages, and thus the Web redirector 320 and redirector servlet 360 are configured to use port number 80 for this HTTP traffic.

Use of the receive channel 330 and send channel 340 to support bi-directional TCP traffic will now be illustrated with reference to an example and with reference to the diagrams in Figs. 4A and 4B. Suppose an end user is using a terminal emulator to communicate with a remote

legacy host application, where the legacy host application is located behind a firewall. The terminal emulator might also reside in a system which is protected by a firewall on the client side. The present invention enables communicating through these firewalls, and through any intermediate HTTP proxies which may be installed between the client device on which the terminal emulator is executing and the target host on which the legacy host application is executing, without requiring change to the client software (i.e. the terminal emulator) or to the remote legacy application software.

The terminal emulator, depicted in the example of Fig. 4A as Telnet client 301, is configured to connect to Web redirector 320 using TCP connection 310. The Telnet client sends TCP messages over this connection as if it is communicating directly to the remote target host, shown in Fig. 4A as Telnet server 381. An example of a client-initiated TCP request message is shown at 400. The Web redirector 320 listens to port 23 for incoming TCP data from client 301. When a TCP message such as message 400 is received from the client, Web redirector 320 inserts that message, in its entirety, into the body of an HTTP POST request message 410. The MIME (Multi-Purpose Internet Mail Extensions) type of this request message is preferably set to “binary/tcp” to prevent it from being misinterpreted by proxies or clients in the network path to redirector servlet 360. References herein to the use of a specific MIME type, however, are for purposes of illustration and not of limitation. The HTTP POST request message 410 is then transmitted from Web redirector 320 to redirector servlet 360 on the send channel 340.

Preferably, the send channel is re-established (by creating a new HTTP POST request) each time a message needs to be sent from a particular client 301 to a particular target server 381.

The redirector servlet 360 receives the inbound HTTP POST message 410 sent from Web redirector 320 on channel 340, and is configured to route these incoming HTTP request messages to the target server (such as server 381). When an incoming HTTP POST request message 410 is received, the redirector servlet 360 extracts the TCP message 400 from the body of this message and forwards the extracted message as a TCP request message 420 over TCP connection 370 to Telnet server 381. To complete the round-trip of the HTTP POST request, an HTTP POST response 430 is then returned from the redirector servlet 360 to the Web redirector 320. The send channel 340 is then preferably closed, and will be re-established when client 301 sends its next TCP request to Web redirector 320.

Preferably, TCP connection 370 is established in response to receiving the first HTTP request message from a particular client which requests port 23 to this target server 381, and this connection is then maintained for the life of the servlet 360. In the preferred embodiment, the first HTTP request message will be an HTTP GET request, sent by Web redirector 320 (as described below with reference to Fig. 4B) to establish a receive channel with the redirector servlet 360. Subsequent incoming HTTP POST and GET messages from the same client to this same target server are preferably associated with this established TCP connection 370 using the session tracking classes and methods which are provided with the Java Web Server Toolkit from Sun Microsystems.

Session tracking classes and methods enable state information to be maintained and made available to servlets. The state information is stored on a server using a session object. This

object is created when a new client session begins, and is kept for the duration of the session. The object stores information about the transactions occurring between the client and the server. (In the present invention, the client for purposes of session tracking is the Web redirector 320 and the server is the redirector servlet 360.) An interface to the object is defined so that servlets can

5 access and modify the state information to reflect the transactions they process for that client. A Session Identifier (“Session ID”) is associated with each client, and is used to correlate a particular session object to the proper client. Preferably, a different session ID is used for each end user client 300 for which Web redirector 320 operates as an HTTP session client. A session ID may be any type of identifier that serves to uniquely identify a particular client to the server.

10 This session ID is then sent as part of the HTTP request syntax for each message sent from the client machine. The server uses the session ID to store information related to the transactions with this client, so that the series of transactions can be treated as a logical on-going communication between the client and the server (instead of simply as random, unrelated messages). The session then encompasses all requests from this client that use this same identifier.

15

20

Session IDs have been implemented on top of the inherently state-less HTTP protocol using two primary approaches. The first is through use of “cookies”. The second is through “URL rewriting”. Cookies and URL rewriting are known in the art, and will not be described in detail herein. The Java Toolkit session tracking classes and methods support both cookies and URL rewriting for passing the session ID between the client and server (and subsequently to servlets).

10
15
20

Fig. 4B illustrates how the present invention supports server-initiated TCP request messages. Suppose Telnet server 381 wishes to send a TCP request to Telnet client 301. Because HTTP is a client-oriented protocol, all messages must originate from the client. The present invention solves this problem by sending an HTTP GET request message (depicted in Fig. 5 4B as HTTP GET request message 460) from the Web redirector 320 on the receive channel 330, where this HTTP GET request serves to open the communications channel through the redirector servlet 360 to the target server 381. As with the other HTTP messages used in the preferred embodiment, the MIME type of HTTP GET request message 460 is preferably set to “binary/tcp”. Web Redirector 320 is always listening on this receive channel 330 for an incoming message from redirector servlet 360. After redirector servlet 360 receives incoming HTTP GET request 460, it performs a TCP read request on the TCP socket for connection 370. This call blocks until the server 381 puts its response onto the TCP socket queue, as in the prior art. When target server 381 has a message to send to client 301, it generates a TCP request 470 and sends that request on TCP connection 370 to redirector servlet 360. Upon receiving TCP request message 470, the redirector servlet’s TCP read request on the socket for connection 370 unblocks. Redirector servlet 360 extracts the server’s TCP request 470 from the body of the HTTP GET response 480, inserts the TCP request 470 into the body of an HTTP GET response message 480, and forwards that HTTP GET response 480 to the Web redirector 320 on receive channel 330. (The MIME type of this HTTP GET response message 480 is also preferably set to 20 “binary/tcp”.) Upon receiving this response 480, thereby completing the round-trip of HTTP GET request 460, the Web redirector 320 (1) sends another HTTP GET request message to redirector servlet 360 (in order to accept the next incoming message generated by target server

381) and (2) extracts the server's TCP request 470 from the body of the HTTP GET response
480. This extracted TCP request is then forwarded to the client 301 on the local TCP connection
310, as shown at 490. This process of listening for incoming messages on the receive channel,
extracting the content, and forwarding it to the client then repeats until the connection between
5 client 301 and server 381 is closed.

As has been demonstrated, the present invention provides a novel technique for enabling
TCP-based messages to be exchanged using HTTP networks and systems, without requiring
change to existing client TCP-based software or to existing TCP-based host systems or
application software. If firewall systems installed in the network between the Web redirector 320
and redirector servlet 360 are configured to encrypt the messages they exchange, then the TCP
messages sent between a particular client and target server are automatically encrypted as well.
16

While the preferred embodiment of the present invention has been described, additional
variations and modifications in that embodiment may occur to those skilled in the art once they
learn of the basic inventive concepts. Therefore, it is intended that the appended claims shall be
15 construed to include both the preferred embodiment and all such variations and modifications as
fall within the spirit and scope of the invention.

What is claimed is:

1. 1. A computer program product for sending Transmission Control Protocol (TCP) messages through HyperText Transfer Protocol (HTTP) systems, the computer program product embodied on one or more computer-readable media and comprising:
 - 4 computer-readable program code means for establishing a send channel from a first component on a client side of a network connection, through one or more HTTP-based systems, to a second component on a remote side of the network connection;
 - 7 computer-readable program code means for establishing a receive channel from the first component, through the one or more HTTP-based systems, to the second component;
 - 9 computer-readable program code means for establishing a first TCP connection from a client on the client side to the first component;
 - 11 computer-readable program code means for establishing a second TCP connection from the second component to a target server on the remote side;
 - 13 computer-readable program code means for transmitting client-initiated TCP requests from the client to the target server on the send channel; and
 - 15 computer-readable program code means for transmitting server-initiated TCP requests from the target server to the client on the receive channel.

1. 2. The computer program product according to Claim 1, wherein the computer-readable program code means for transmitting client-initiated TCP requests further comprises:
 - 3 computer-readable program code means for receiving a client-initiated TCP request from the client at the first component on the first TCP connection;

5 computer-readable program code means for packaging the received client-initiated TCP
6 request in an HTTP POST request message;
7 computer-readable program code means for sending the HTTP POST request message to
8 the second component on the network connection;
9 computer-readable program code means for receiving the sent HTTP POST request
10 message at the second component;
11 computer-readable program code means for extracting the client-initiated TCP request
12 from the received HTTP POST request message; and
13 computer-readable program code means for forwarding the extracted client-initiated TCP
14 request to the target server on the second TCP connection.

3. The computer program product according to Claim 2, wherein the computer-readable
program code means for transmitting client-initiated TCP requests further comprises computer-
readable program code means for acknowledging the HTTP POST request by sending an HTTP
POST response from the second component to the first component on the network connection.

1 4. The computer program product according to Claim 3, wherein the computer-readable
2 program code means for establishing the send channel operates in response to the computer-
3 readable program code means for receiving the client-initiated TCP request, and wherein the
4 computer-readable program code means for transmitting client-initiated TCP requests further
5 comprises:
6 computer-readable program code means for receiving the HTTP POST response at the

7 first component; and

8 computer-readable program code means for closing the send channel, responsive to
9 operation of the computer-readable program code means for receiving the HTTP POST response.

1 5. The computer program product according to Claim 1, wherein the computer-readable
2 program code means for transmitting server-initiated TCP requests further comprises:

3 computer-readable program code means for sending an HTTP GET request message from
4 the first component to the second component on the network connection;

5 computer-readable program code means for receiving the sent HTTP GET request
6 message at the second component;

7 computer-readable program code means for receiving a server-initiated TCP request from
8 the target server at the second component on the second TCP connection;

9 computer-readable program code means for packaging the received server-initiated TCP
10 request in an HTTP GET response message which acknowledges the received HTTP GET
11 request message;

12 computer-readable program code means for sending the HTTP GET response message
13 from the second component to the first component on the network connection;

14 computer-readable program code means for receiving the sent HTTP GET response
15 message at the first component;

16 computer-readable program code means for extracting the server-initiated TCP request
17 from the received HTTP GET response message; and

18 computer-readable program code means for forwarding the extracted server-initiated TCP

19 request to the client on the first TCP connection.

1 6. The computer program product according to Claim 5, wherein the computer-readable
2 program code means for transmitting server-initiated TCP requests further comprises:

3 computer-readable program code means for performing a read operation on the second
4 TCP connection, responsive to operation of the computer-readable program code means for
5 receiving the sent HTTP GET request message and prior to operation of the computer-readable
6 program code means for receiving the server-initiated TCP request; and

7 computer-readable program code means for using the received server-initiated TCP
8 request as a result of the read operation, thereby triggering operation of the computer-readable
9 program code means for packaging the received server-initiated TCP request in the HTTP GET
10 response message.

1 7. The computer program product according to Claim 5, wherein the computer-readable
2 program code means for transmitting server-initiated TCP requests further comprises computer-
3 readable program code means for preparing to receive another server-initiated TCP request by
4 triggering operation of the computer-readable program code means for sending the HTTP GET
5 request message from the first component to the second component, responsive to operation of
6 the computer-readable program code means for receiving the sent HTTP GET response message
7 at the first component.

1 8. The computer program product according to Claim 2, wherein a Multi-Purpose Internet

2 Mail Extensions (MIME) type of the HTTP POST request message is set to “binary/tcp”.

1 9. The computer program product according to Claim 5, wherein a Multi-Purpose Internet
2 Mail Extensions (MIME) type of the HTTP GET request message is set to “binary/tcp”.

1 10. A system for sending Transmission Control Protocol (TCP) messages through HyperText
2 Transfer Protocol (HTTP) systems, comprising:

3 means for establishing a send channel from a first component on a client side of a network
4 connection, through one or more HTTP-based systems, to a second component on a remote side
5 of the network connection;

6 means for establishing a receive channel from the first component, through the one or
7 more HTTP-based systems, to the second component;

8 means for establishing a first TCP connection from a client on the client side to the first
9 component;

10 means for establishing a second TCP connection from the second component to a target
11 server on the remote side;

12 means for transmitting client-initiated TCP requests from the client to the target server on
13 the send channel; and

14 means for transmitting server-initiated TCP requests from the target server to the client on
15 the receive channel.

1 11. The system according to Claim 10, wherein the means for transmitting client-initiated TCP

2 requests further comprises:

3 means for receiving a client-initiated TCP request from the client at the first component on
4 the first TCP connection;

5 means for packaging the received client-initiated TCP request in an HTTP POST request
6 message;

7 means for sending the HTTP POST request message to the second component on the
8 network connection;

9 means for receiving the sent HTTP POST request message at the second component;

10 means for extracting the client-initiated TCP request from the received HTTP POST
11 request message; and

12 means for forwarding the extracted client-initiated TCP request to the target server on the
13 second TCP connection.

12. The system according to Claim 11, wherein the means for transmitting client-initiated TCP
requests further comprises means for acknowledging the HTTP POST request by sending an
HTTP POST response from the second component to the first component on the network
connection.

1 13. The system according to Claim 12, wherein the means for establishing the send channel
2 operates in response to the means for receiving the client-initiated TCP request, and wherein the
3 means for transmitting client-initiated TCP requests further comprises:

4 means for receiving the HTTP POST response at the first component; and

5 means for closing the send channel, responsive to operation of the means for receiving the
6 HTTP POST response.

1 14. The system according to Claim 10, wherein the means for transmitting server-initiated
2 TCP requests further comprises:

3 means for sending an HTTP GET request message from the first component to the second
4 component on the network connection;

5 means for receiving the sent HTTP GET request message at the second component;

6 means for receiving a server-initiated TCP request from the target server at the second
7 component on the second TCP connection;

8 means for packaging the received server-initiated TCP request in an HTTP GET response
9 message which acknowledges the received HTTP GET request message;

10 means for sending the HTTP GET response message from the second component to the
11 first component on the network connection;

12 means for receiving the sent HTTP GET response message at the first component;

13 means for extracting the server-initiated TCP request from the received HTTP GET
14 response message; and

15 means for forwarding the extracted server-initiated TCP request to the client on the first
16 TCP connection.

1 15. The system according to Claim 14, wherein the means for transmitting server-initiated
2 TCP requests further comprises:

3 means for performing a read operation on the second TCP connection, responsive to
4 operation of the means for receiving the sent HTTP GET request message and prior to operation
5 of the means for receiving the server-initiated TCP request; and
6 means for using the received server-initiated TCP request as a result of the read operation,
7 thereby triggering operation of the means for packaging the received server-initiated TCP request
8 in the HTTP GET response message.

1 16. The system according to Claim 14, wherein the means for transmitting server-initiated
2 TCP requests further comprises means for preparing to receive another server-initiated TCP
3 request by triggering operation of the means for sending the HTTP GET request message from
4 the first component to the second component, responsive to operation of the means for receiving
5 the sent HTTP GET response message at the first component.

1 17. The system according to Claim 11, wherein a Multi-Purpose Internet Mail Extensions
2 (MIME) type of the HTTP POST request message is set to “binary/tcp”.

1 18. The system according to Claim 14, wherein a Multi-Purpose Internet Mail Extensions
2 (MIME) type of the HTTP GET request message is set to “binary/tcp”.

1 19. A method for sending Transmission Control Protocol (TCP) messages through HyperText
2 Transfer Protocol (HTTP) systems, comprising the steps of:
3 establishing a send channel from a first component on a client side of a network

4 connection, through one or more HTTP-based systems, to a second component on a remote side
5 of the network connection;

6 establishing a receive channel from the first component, through the one or more HTTP-
7 based systems, to the second component;

8 establishing a first TCP connection from a client on the client side to the first component;

9 establishing a second TCP connection from the second component to a target server on
10 the remote side;

11 transmitting client-initiated TCP requests from the client to the target server on the send
12 channel; and

13 transmitting server-initiated TCP requests from the target server to the client on the
14 receive channel.

20. The method according to Claim 19, wherein the step of transmitting client-initiated TCP
requests further comprises the steps of:

1 receiving a client-initiated TCP request from the client at the first component on the first
2 TCP connection;

3 packaging the received client-initiated TCP request in an HTTP POST request message;

4 sending the HTTP POST request message to the second component on the network
5 connection;

6 receiving the sent HTTP POST request message at the second component;

7 extracting the client-initiated TCP request from the received HTTP POST request
8 message; and

11 forwarding the extracted client-initiated TCP request to the target server on the second
12 TCP connection.

1 21. The method according to Claim 20, wherein the step of transmitting client-initiated TCP
2 requests further comprises the step of acknowledging the HTTP POST request by sending an
3 HTTP POST response from the second component to the first component on the network
4 connection.

1 22. The method according to Claim 21, wherein the step of establishing the send channel
2 operates in response to the step of receiving the client-initiated TCP request, and wherein the step
3 of transmitting client-initiated TCP requests further comprises the steps of:
4 receiving the HTTP POST response at the first component; and
5 closing the send channel, responsive to receiving the HTTP POST response.

1 23. The method according to Claim 19, wherein the step of transmitting server-initiated TCP
2 requests further comprises the steps of:
3 sending an HTTP GET request message from the first component to the second
4 component on the network connection;
5 receiving the sent HTTP GET request message at the second component;
6 receiving a server-initiated TCP request from the target server at the second component
7 on the second TCP connection;
8 packaging the received server-initiated TCP request in an HTTP GET response message

9 which acknowledges the received HTTP GET request message;

10 sending the HTTP GET response message from the second component to the first

11 component on the network connection;

12 receiving the sent HTTP GET response message at the first component;

13 extracting the server-initiated TCP request from the received HTTP GET response

14 message; and

15 forwarding the extracted server-initiated TCP request to the client on the first TCP

16 connection.

24. The method according to Claim 23, wherein the step of transmitting server-initiated TCP
requests further comprises the steps of:

1 performing a read operation on the second TCP connection, responsive to receiving the
2 sent HTTP GET request message and prior to receiving the server-initiated TCP request; and
3 using the received server-initiated TCP request as a result of the read operation, thereby
4 triggering the step of packaging the received server-initiated TCP request in the HTTP GET
5 response message.

1 25. The method according to Claim 23, wherein the step of transmitting server-initiated TCP
2 requests further comprises the step of preparing to receive another server-initiated TCP request
3 by triggering the step of sending the HTTP GET request message from the first component to the
4 second component, responsive to receiving the sent HTTP GET response message at the first
5 component.

1 26. The method according to Claim 20, wherein a Multi-Purpose Internet Mail Extensions
2 (MIME) type of the HTTP POST request message is set to “binary/tcp”.

1 27. The method according to Claim 23, wherein a Multi-Purpose Internet Mail Extensions
2 (MIME) type of the HTTP GET request message is set to “binary/tcp”.

1 28. A method for transporting bi-directional protocol traffic through uni-directional protocol
2 systems, comprising the steps of:

establishing a send channel from a first component on a client side of a network connection, through one or more uni-directional protocol-based systems, to a second component on a remote side of the network connection;

establishing a receive channel from the first component, through the one or more unidirectional protocol-based systems, to the second component;

establishing a first bi-directional protocol connection from a client on the client side to the first component;

10 establishing a second bi-directional protocol connection from the second component to a
11 target server on the remote side;

12 transmitting client-initiated bi-directional protocol requests from the client to the target
13 server on the send channel; and

transmitting server-initiated bi-directional protocol requests from the target server to the client on the receive channel.

1 29. The method according to Claim 28, wherein the step of transmitting client-initiated bi-
2 directional protocol requests further comprises the steps of:

3 receiving a client-initiated bi-directional protocol request from the client at the first
4 component on the first bi-directional protocol connection;
5 packaging the received client-initiated bi-directional protocol request in a uni-directional
6 protocol write request message;
7 sending the uni-directional protocol write request message to the second component on
8 the network connection;

9 receiving the sent uni-directional protocol write request message at the second
10 component;

11 extracting the client-initiated bi-directional protocol request from the received uni-
12 directional protocol write request message; and
13 forwarding the extracted client-initiated bi-directional protocol request to the target server
14 on the second bi-directional protocol connection.

1 30. The method according to Claim 28, wherein the step of transmitting server-initiated bi-
2 directional protocol requests further comprises the steps of:

3 sending a uni-directional protocol read request message from the first component to the
4 second component on the network connection;
5 receiving the sent uni-directional protocol read request message at the second component;
6 receiving a server-initiated bi-directional protocol request from the target server at the

7 second component on the second bi-directional protocol connection;

8 packaging the received server-initiated bi-directional protocol request in a uni-directional

9 protocol read response message which acknowledges the received uni-directional protocol read

10 request message;

11 sending the uni-directional protocol read response message from the second component to

12 the first component on the network connection;

13 receiving the sent uni-directional protocol read response message at the first component;

14 extracting the server-initiated bi-directional protocol request from the received uni-

15 directional protocol read response message; and

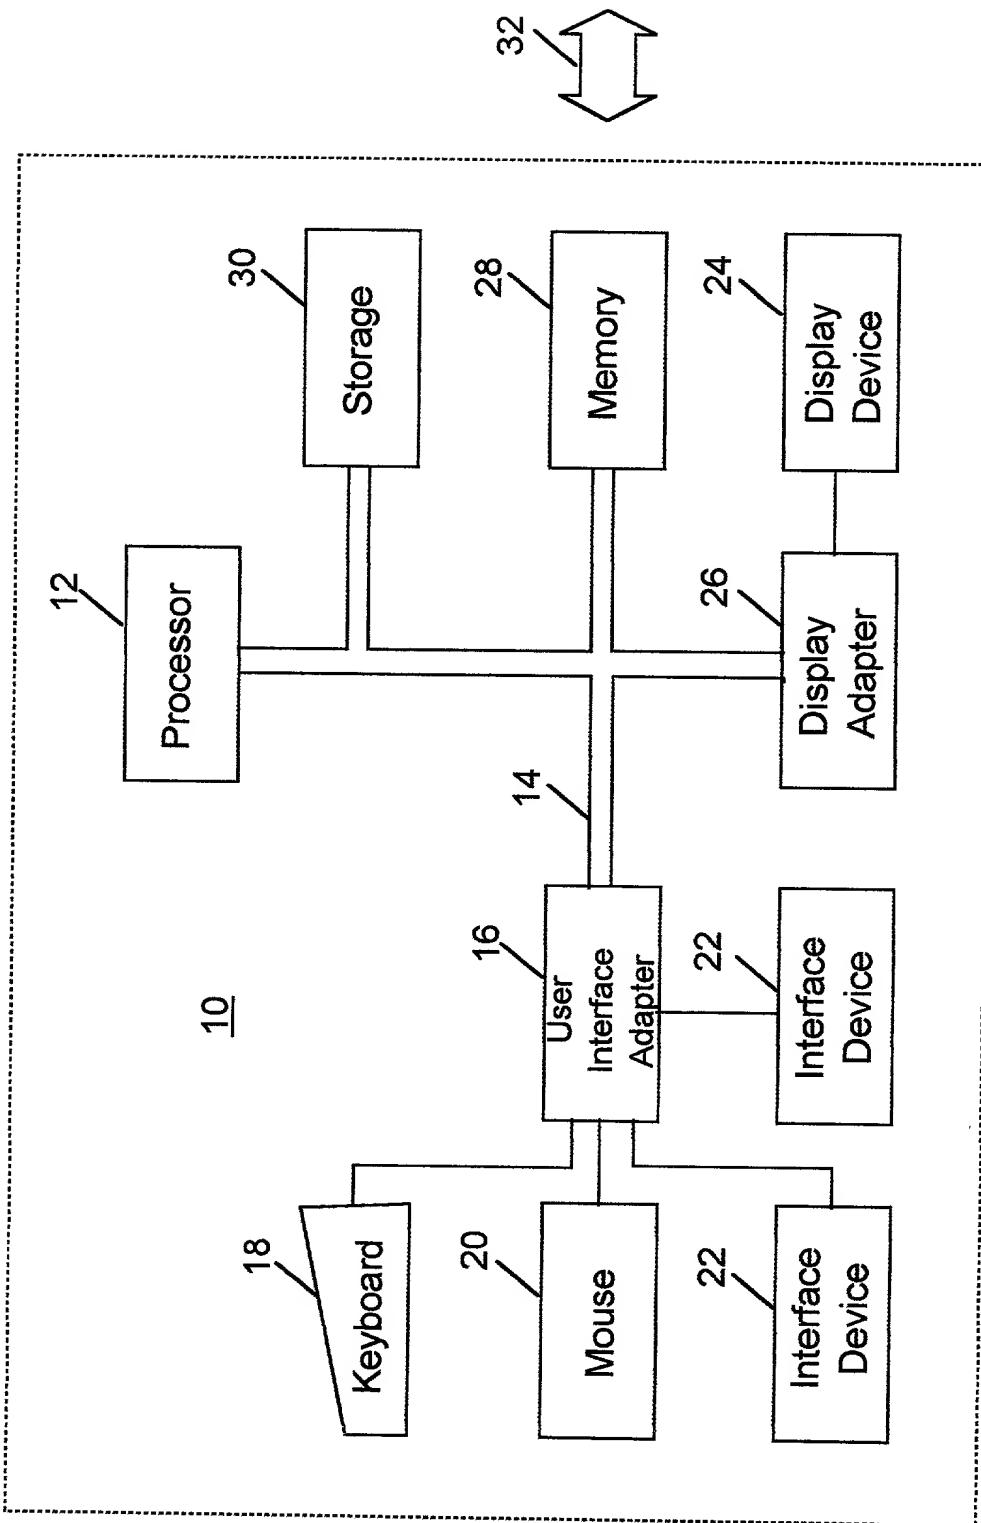
16 forwarding the extracted server-initiated bi-directional protocol request to the client on the

17 first bi-directional protocol connection.

ABSTRACT

The present invention provides a method, system, and computer program product for facilitating end-user access to remote applications based on TCP (Transmission Control Protocol). TCP traffic is packaged into the body of HTTP (HyperText Transfer Protocol) messages, enabling the TCP traffic to be transmitted through HTTP systems in an application-independent manner as if it were native HTTP traffic. TCP traffic therefore flows through firewall systems without requiring access to an external TCP port, and flows through HTTP proxy servers as well. The disclosed technique enables mapping bi-directional TCP messages to HTTP, which is designed for client-initiated message traffic only.

**FIG. 1
(PRIOR ART)**



**FIG. 2
(PRIOR ART)**

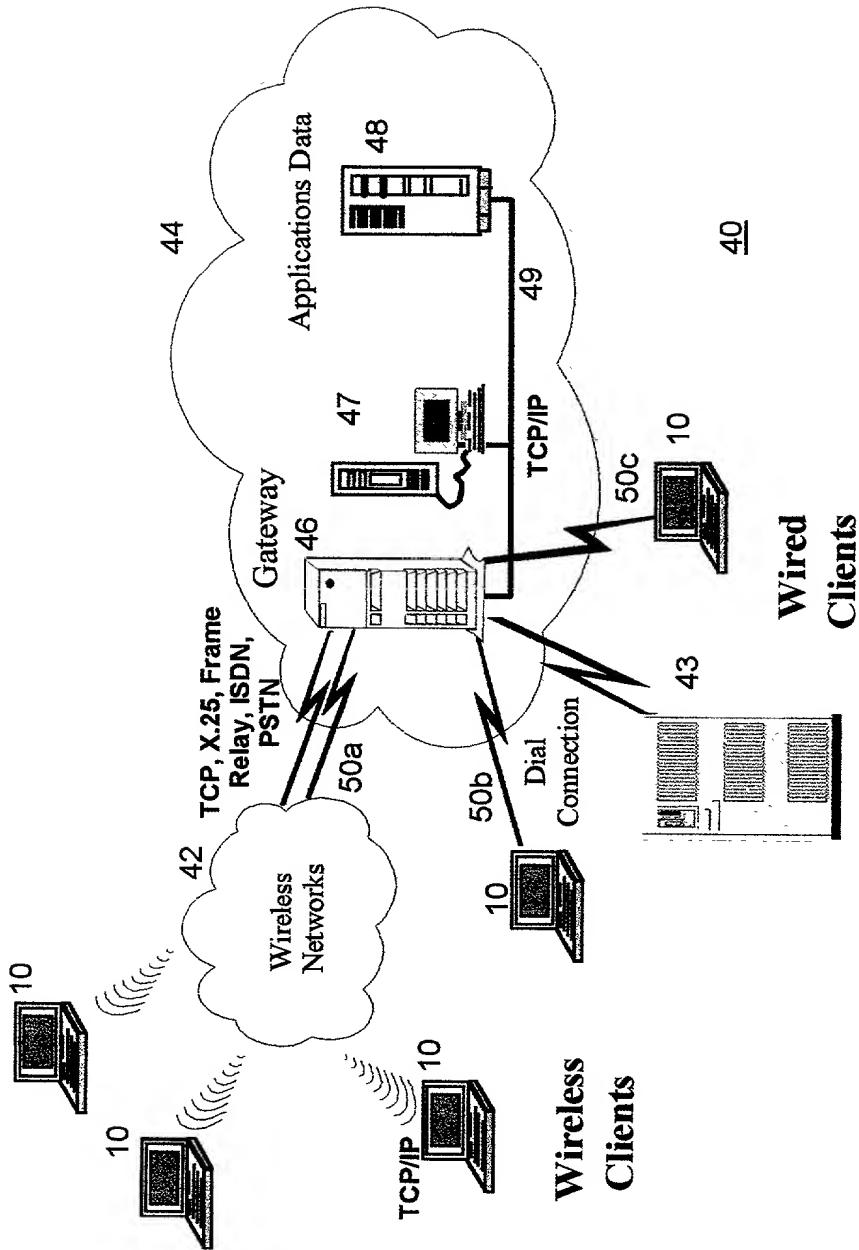


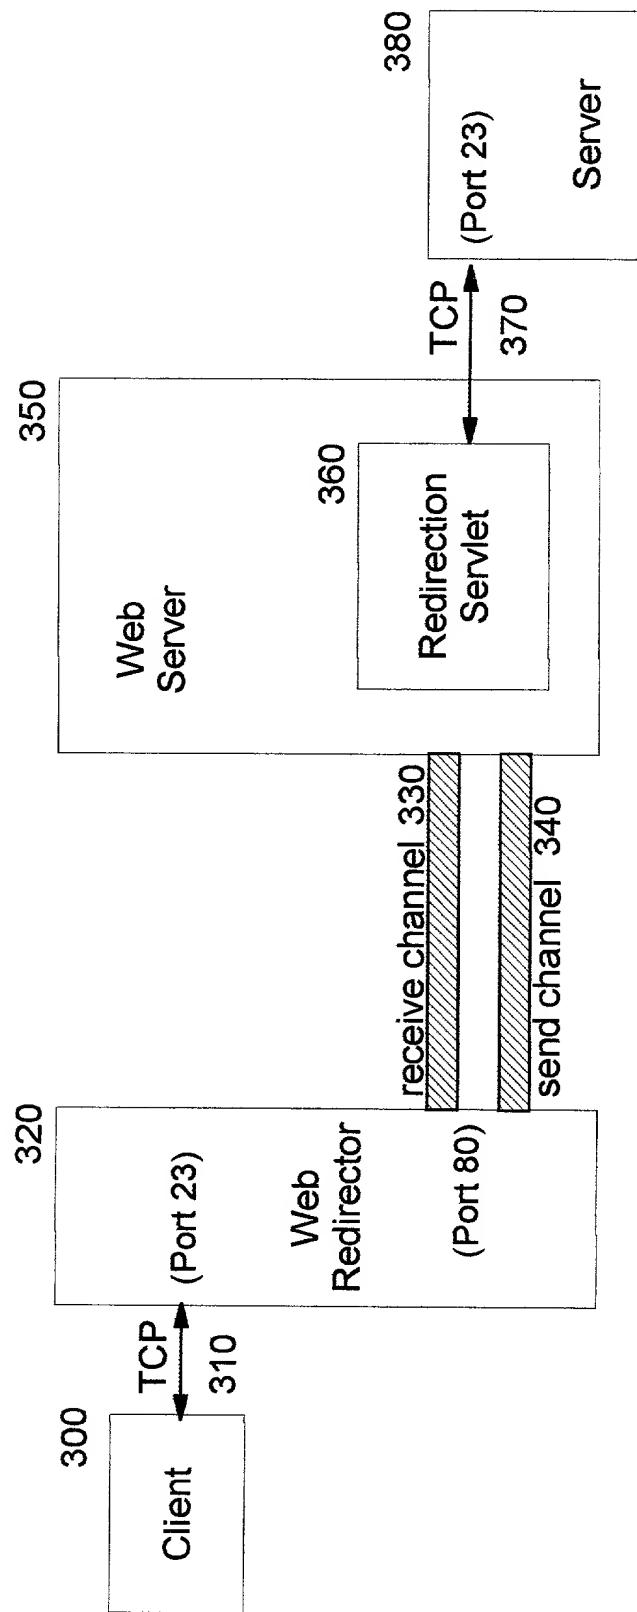
FIG. 3

FIG. 4A

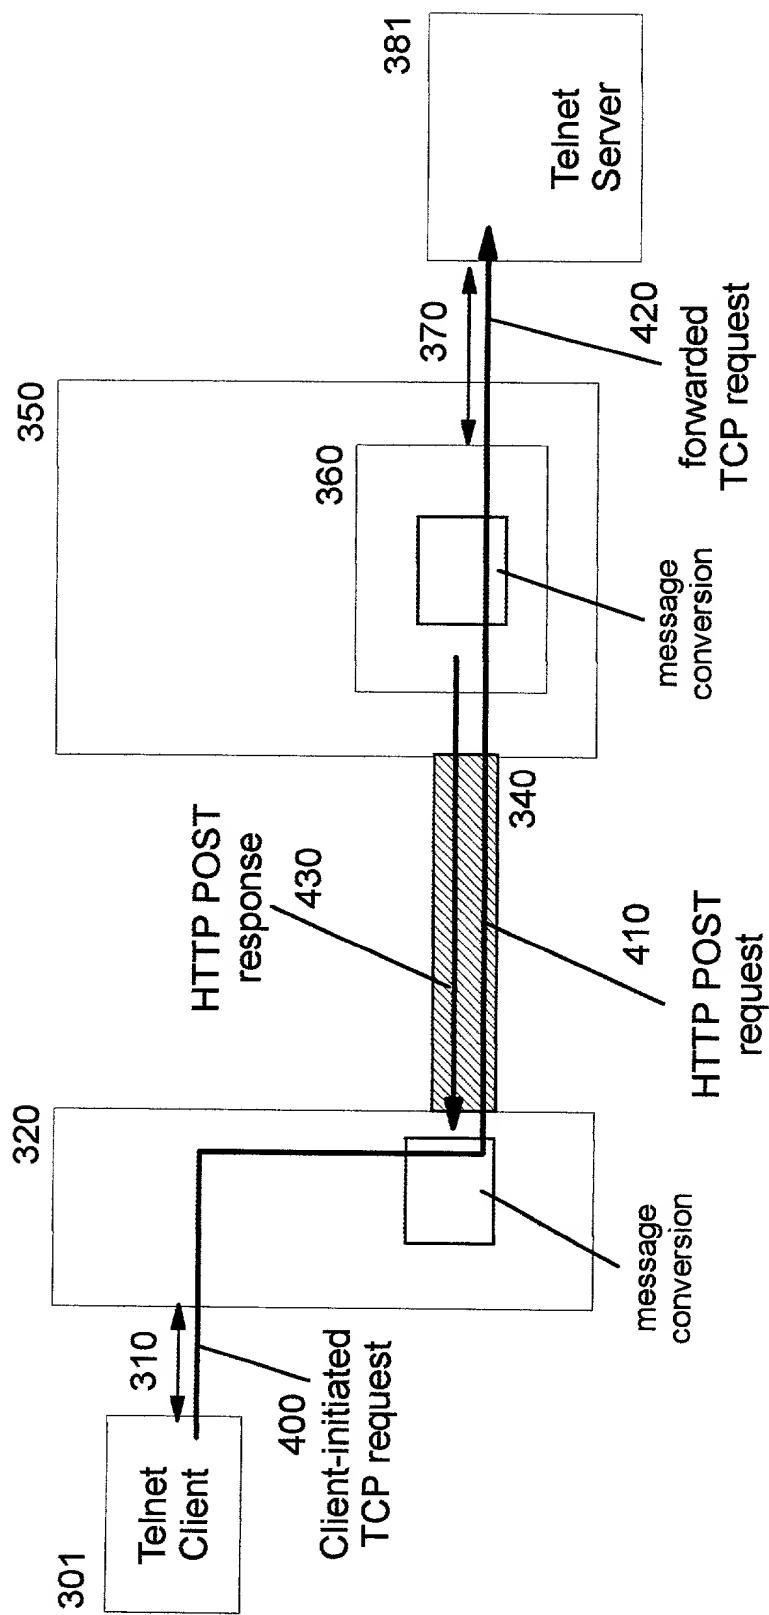
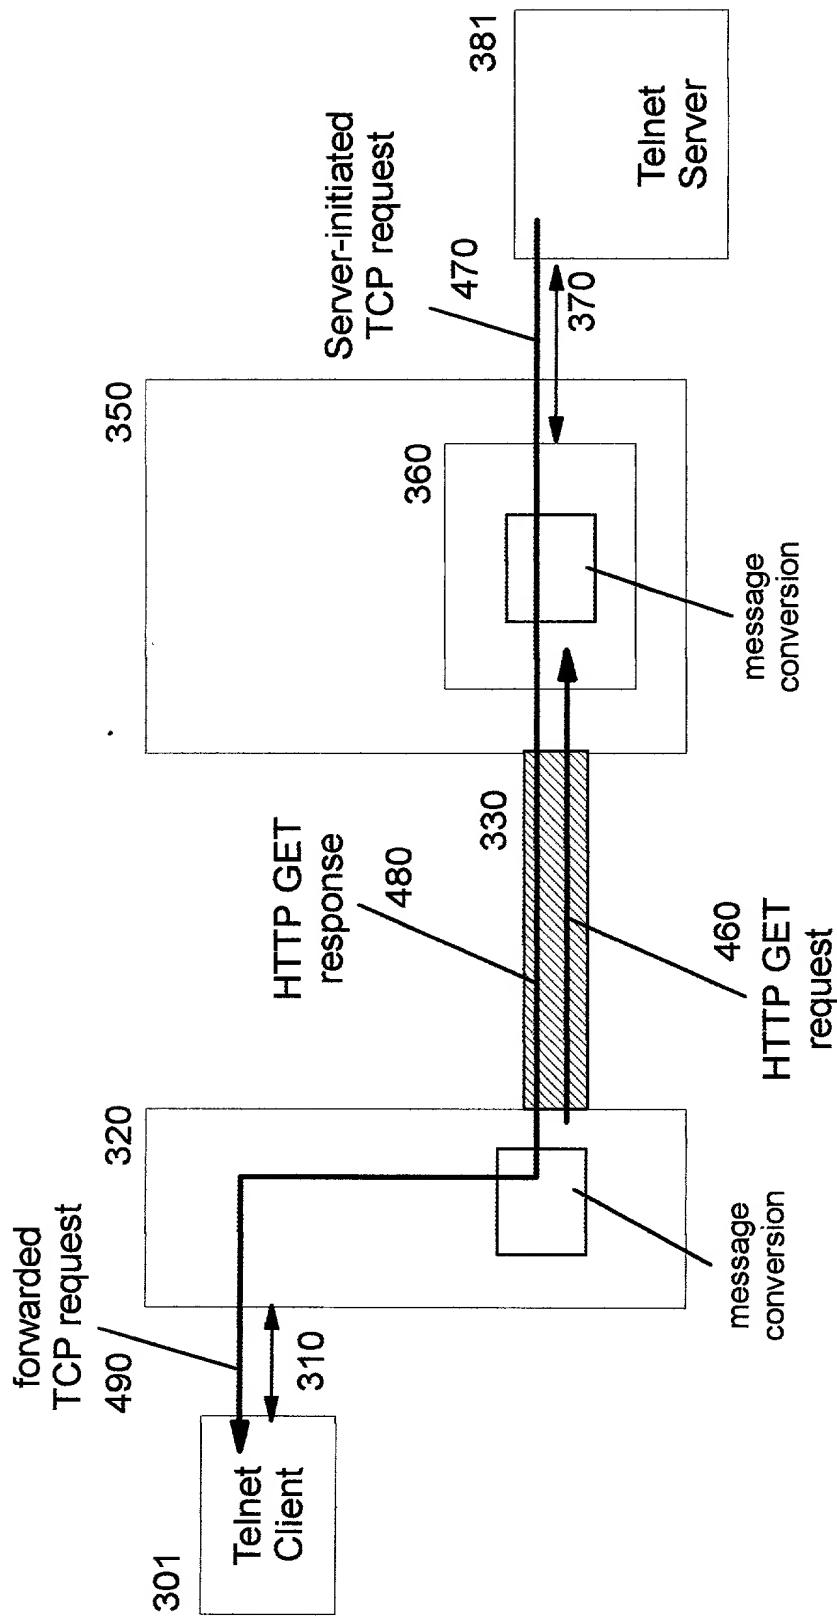


FIG. 4B



**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION**

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name; I believe I am an original, first and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled:

Technique for Sending TCP Messages through HTTP Systems

the specification of which is identified by the attorney (IBM) Docket Number appearing above.

I hereby state that I have reviewed and understand the contents of the above- identified specification, including the claims.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

<u>Number</u>	<u>Country</u>	<u>Day/Month/Year</u>	<u>Priority Claimed</u>
---------------	----------------	-----------------------	-------------------------

I hereby claim the benefit (a) under Title 35, United States Code, §119(e) of any U.S. application listed below and identified as a provisional application or (b) under Title 35, United States Code, §120 of any U.S. application listed below and not identified as a provisional application, and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior U.S. application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in Title 37, Code of Federal Regulations, §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

Prior U.S. Applications

<u>Serial No.</u>	<u>Filing Date</u>	<u>Status</u>
-------------------	--------------------	---------------

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Title 18, United States Code, §1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Bruce A. Clay, Reg. No. 32,121; Gregory M. Doudnikoff, Reg. No. 32,847; Edward H. Duffield, Reg. No. 25,970; Jerry W. Herndon, Reg. No. 27,901; Gerald R. Woods, Reg. No. 24,144; Jeanine S. Ray-Yarletts, Reg. No. 39,808; Joseph C. Redmond, Jr., Reg. No. 18,753; John E. Hoel, Reg. No. 26,279; Christopher A. Hughes, Reg. No. 26,914; and Edward A. Pennington, Reg. No. 32,588;

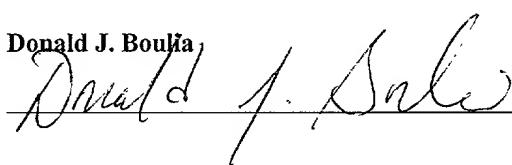
AND also, Marcia L. Doubet, Reg. No. 40,999

Send all correspondence to: Jeanine S. Ray-Yarletts, IBM Corporation 972/B656; PO Box 12195; Research Triangle Park, NC 27709.

Inventor:

Donald J. Boula

Signature:



7/14/00

Date

Residence: 304 Burgwin Wright Way, Apex, North Carolina 27502

Citizenship: United States of America

Post Office Address: same as residence